

An algorithm for improving Non-Local Means operators via low-rank approximation

Victor May^{* 1}, Yosi Keller^{† 2}, Nir Sharon^{‡1}, and Yoel Shkolnisky^{§ 1}

¹School of Mathematical Sciences, Tel-Aviv University, Tel-Aviv, Israel

²Faculty of Engineering, Bar-Ilan University, Ramat-Gan, Israel

Abstract

We present a method for improving a Non Local Means operator by computing its low-rank approximation. The low-rank operator is constructed by applying a filter to the spectrum of the original Non Local Means operator. This results in an operator which is less sensitive to noise while preserving important properties of the original operator. The method is efficiently implemented based on Chebyshev polynomials and is demonstrated on the application of natural images denoising. For this application, we provide a comprehensive comparison of our method with leading denoising methods.

Keywords: Denoising, Non-Local Means operator, Chebyshev polynomials.

1 Introduction

Denoising images is a classical problem in image processing. Representative approaches for solving this problem include local methods such as linear filtering and anisotropic diffusion [13], global methods such as total-variation [14] and wavelet shrinkage [3], and discrete methods such as Markov Random field denoising [16] and discrete universal denoiser [12].

As is often the case with inverse problems, many of the algorithms for image denoising involve priors on the solution. Commonly, the only prior knowledge assumed about the

^{*}mayvic@gmail.com

[†]yosi.keller@gmail.com

[‡]nir.sharon@math.tau.ac.il

[§]yoelsh@post.tau.ac.il

image is that it is of natural origin. In that case, the sought-after prior should represent the statistics of a natural image. Natural image statistics is a research topic on its own [10, 20], having importance for image processing problems other than denoising: segmentation, texture synthesis, image inpainting, super-resolution and more. An important observation in natural image statistics is that images contain repetitive local patterns. This observation is the basis of patch-based denoising methods such as Non-Local Means (NLM) [1] and block-matching and 3D filtering (BM3D) [4].

In NLM, the denoising operator is constructed from patches of the corrupted image itself. As such, the denoising operator is affected by the noise. It had been noted [11, 15] that larger eigenvalues of the operator are less sensitive to noise than smaller ones, creating an opportunity to improve the operator. In this work we propose a method for computing such a modified operator, by means of filtering the noisy one with a low-pass filter applied to its eigenvalues. In other words, we pose the problem as a filter design task, which is a classical task in signal processing. Our chosen filter function suppresses eigenvalues with small magnitude, and accurately approximates the remaining low-rank operator, while preserving the fundamental properties of the original operator. The filter is efficiently applied to the NLM operator based on Chebyshev polynomials for matrices.

We further study the concept of low-rank NLM denoising operators by numerical experiments. In these experiments we investigate our method for two main questions. The first is the dependence of the low-rank operator on its two tuning parameters, which are the width of the kernel of the original NLM operator, and the rank of the low-rank operator. The second question is the performance of our method compared to other advanced denoising algorithms. We provide a comprehensive comparison which includes a few popular state-of-the-art methods and a two-stage denoising scheme based on our low-rank operator.

The paper is organized as follows. In Section 2 we present the notations, the problem’s formulation and the NLM method. In Section 3 we introduce our method of constructing low-rank approximations for NLM operators, which is based on matrix filtering functions. Next, in Section 4, we show how to efficiently implement this low-rank approximation using Chebyshev polynomials. Section 5 provides numerical experiments with our algorithm, where we discuss its parameters and compare it to other advanced denoising methods. We summarize the paper in Section 6.

2 Preliminaries

We start by introducing the notations and the model for the denoising problem. Let I be a discrete grid of dimension d (typically, $d = 1$ or $d = 2$) and denote by $X = \{x_i \mid i \in I\}$ a

clean signal defined on I . Let $Y = \{y_i \mid i \in I\}$ be a signal contaminated with noise, that is $y_i = x_i + r_i$, where $r_i \sim N(0, \sigma^2)$, $i \in I$, are independently identically distributed Gaussian random variables. The goal is to estimate X given Y .

The non-local means (NLM) estimator is given as follows. Denote by $N_i^Y = N_i^Y(d, p)$ the indices of the pixels within a hypercube of dimension d and side length p centered around the pixel $y_i \in Y$. Let $v(N_i^Y)$, $i \in I$ be the values of the pixels of Y at the indices N_i^Y , treated as a column vector. The NLM algorithm estimates x_i as

$$\hat{x}_i = \frac{1}{z_i} \sum_{y_j \in Y} K_h \left(v(N_i^Y) - v(N_j^Y) \right) y_j,$$

where $K_h(\cdot) = \exp(-\frac{\|\cdot\|_2^2}{2h^2})$ is the Gaussian kernel function with width $h > 0$, and

$$z_i = \sum_{y_j \in Y} K_h \left(v(N_i^Y) - v(N_j^Y) \right)$$

is a normalization factor.

In matrix notation the NLM estimator is written as

$$\hat{x} = Ay, \tag{1}$$

where $y \in \mathbb{R}^{n \times 1}$ is the original noisy signal Y represented as a vector, \hat{x} is the denoised signal (the estimations vector), and A is a NLM operator of Y , defined as follows.

Definition 1. A matrix $A \in \mathbb{R}^{n \times n}$ is a NLM operator of Y if $A = D^{-1}W$, where

$$W_{ij} = K_h \left(v(N_i^Y) - v(N_j^Y) \right)$$

with $K_h(\cdot) = \exp(-\frac{\|\cdot\|_2^2}{2h^2})$ and $h > 0$, and D is a diagonal matrix given by $D_{ii} = \sum_{j=1}^n W_{ij}$.

Remark 1. The pseudocode of all algorithms described in this paper is given in Appendix B. In these algorithms, we denote by $\text{NLM}_{p,h}(Y)$ the NLM operator of image Y with patch size p and kernel width h .

3 Denoising the NLM operator

As the NLM operator A in (1) is constructed from noisy data, we would like to replace it with an operator which is less noisy. One way is to replace A with its low-rank approximation

[11, 15]. A low-rank approximation obtained by a truncated singular-value decomposition (SVD) is optimal under the L_2 norm [8, Chapter 2.5], however, directly computing the SVD of A is computationally expensive and often impractical. Computing a truncated eigenvalues (spectral) decomposition is an alternative method, typically implemented by a variant of the power iterations algorithm, such as Lanczos iterations. Unfortunately, all existing spectral decomposition methods are often too slow to be applied to a NLM matrix. Low rank approximation of the NLM matrix via truncated spectral decomposition is given in Algorithm 1 in Appendix B.

Our approach is to construct a low-rank approximation of a NLM operator by using a matrix function $f: \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times n}$, and computing $f(A)$, where A is a NLM operator, as in Definition 1. We next characterize functions of NLM operators and suggest properties of such functions which are suitable for our purposes. Then, we present a particular family of functions with controlled low-rank that have these properties.

3.1 Matrix functions of NLM operators

Before studying matrix functions on NLM operators, we summarize a few properties of these NLM matrices.

Lemma 2. *Let A be a NLM operator, as defined in Definition 1. Then A has the following properties:*

1. *A is positive diagonalizable, namely there exists an invertible matrix Q that satisfies $A = Q^{-1}\Lambda Q$, where Λ is a diagonal matrix $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$, and $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n \geq 0$.*
2. *$\lambda_1 = 1$ is the maximal eigenvalue of A with a corresponding eigenvector $\mathbf{1}$ (the all-ones vector), that is, $A\mathbf{1} = \mathbf{1}$.*

The proof is given in Appendix A. As a first conclusion from Lemma 2 we have

Corollary 3. *In the notation of Lemma 2, any matrix function $f: \mathbb{R}^{n \times n} \rightarrow \mathbb{R}^{n \times n}$ satisfies*

$$f(A) = Q^{-1}f(\Lambda)Q.$$

The proof follows directly from the diagonalizability of A and the definition of matrix functions, see e.g., [8, Corollary 11.1.2].

We design f so that $f(A)$ is a low-rank approximation of the NLM operator A . Moreover, we wish to guarantee that $f(A)$ retains the properties of A listed in Lemma 2. The first property in Lemma 2 follows directly from Corollary 3 and allows for repeated applications

of the operator $f(A)$, as the eigenvalues are non-negative and bounded by 1. The advantages of repeated applications of a denoising operator have been demonstrated in [15]. The second property in Lemma 2, that is $f(A)\mathbf{1} = \mathbf{1}$ (row stochastic normalization of A), is useful because then the elements of $y = f(A)x$ are affine combinations (linear combinations whose coefficients sum to one) of the elements of x . In particular it is desired for a denoising operator to map a constant signal to itself, see e.g., [3].

Summarizing the above discussion of desired properties of the matrix $f(A)$, we define the notion of an extended NLM operator.

Definition 4. *A matrix $B \in \mathbb{R}^{n \times n}$ is an extended NLM operator if B is diagonalizable, such that $B\mathbf{1} = \mathbf{1}$, with all its eigenvalues in $[0, 1]$.*

While it is possible to define an extended NLM operator such that its eigenvalues are in $(-1, 1]$, we prefer for simplicity to use Definition 4 above, as this allows for a richer family of functions to be used in Section 3.2 below. This is consistent with the Gaussian kernel used in Definition 1, which ensures that the resulting operator is non-negative definite.

Equipped with the latter we have,

Theorem 5. *Let A be a NLM operator. The following conditions on f are sufficient for $f(A)$ to be an extended NLM operator.*

1. $f(1) = 1$.
2. f is defined on the interval $[0, 1]$ and $\{f(x) \mid 0 \leq x \leq 1\} \subseteq [0, 1]$.

Proof. By Corollary 3, f acts solely on the diagonal matrix Λ , meaning that

$$f(A) = Q^{-1} \text{diag}(f(\lambda_1), \dots, f(\lambda_n)) Q,$$

where the eigenvalues of A are ordered such that $\lambda_1 = 1$. Therefore, the second condition guarantees that the eigenvalues of $f(A)$ are in $[0, 1]$. In addition, the first condition ensures the maximal eigenvalue is indeed one, and is given by $f(\lambda_1)$, that is $f(\lambda_1) = 1$. The eigenvector corresponding to $f(\lambda_1)$ is the first column of Q , which is $\mathbf{1}$ (up to a scalar scale), namely, $f(A)\mathbf{1} = \mathbf{1}$. \square

The following corollary allows composition of functions satisfying the conditions of Theorem 5.

Corollary 6. *Let B be an extended NLM operator, and let g be a function satisfying the conditions of Theorem 5. Then, $g(B)$ is also an extended NLM operator.*

The proof is elementary and thus omitted.

3.2 Constructing a low-rank extended NLM operator

We would like to construct a function $f(A)$ that satisfies the two conditions from Theorem 5, and moreover, suppresses the noise in the NLM operator A . A natural option is to choose a function that retains the eigenvalues above a given threshold. By Corollary 3, choosing such a function reduces to choosing an appropriate scalar function.

Our first prototype for a scalar thresholding function is

$$g_\omega(x) = \begin{cases} 0 & x < \omega, \\ x & \omega \leq x. \end{cases} \quad (2)$$

This function satisfies the conditions of Theorem 5 for $0 \leq \omega \leq 1$, while zeroing values lower than ω and acting as the identity for higher values. However, this function is not smooth which eventually results in its slow evaluation for matrices (a detailed discussion is given in Section 4.2).

Inspired by the gain function of the Butterworth filter [2] (also known as maximal flat filter)

$$f_{\omega,d}^b(x) = \left(1 + \left(\frac{1-x}{1-\omega}\right)^{2d}\right)^{-\frac{1}{2}}, \quad x \in [0, 1],$$

which approximates the ω -shifted Heaveside function on the segment $[0, 1]$, we propose to use

$$f_{\omega,d}^{sb}(x) = x \left(1 + \left(\frac{1-x}{1-\omega}\right)^{2d}\right)^{-\frac{1}{2}}, \quad x \in [0, 1]. \quad (3)$$

We term this function the Slanted Butterworth (SB) function. The SB function serves as an approximation to g_ω of (2). The response of both functions $f_{\omega,d}^b(x)$ and $f_{\omega,d}^{sb}(x)$ for various values of d is shown in Figure 1.

In its matrix version, the SB function becomes

$$f_{\omega,d}^{sb}(A) = A \left(I + \left(\frac{1}{1-\omega} (I - A) \right)^{2d} \right)^{-\frac{1}{2}}, \quad (4)$$

where $A \in \mathbb{R}^{n \times n}$ and I is the $n \times n$ identity matrix. Combining Theorem 5 and Corollary 6, one can verify that $f_{\omega,d}^{sb}(A)$, $0 \leq \omega \leq 1$, is indeed an extended NLM operator and thus $f_{\omega,d}^{sb}$ is applicable for our purposes. The parameter $0 \leq \omega \leq 1$ is termed the filter cutoff and $d \in \mathbb{N}$ is the order of the filter.

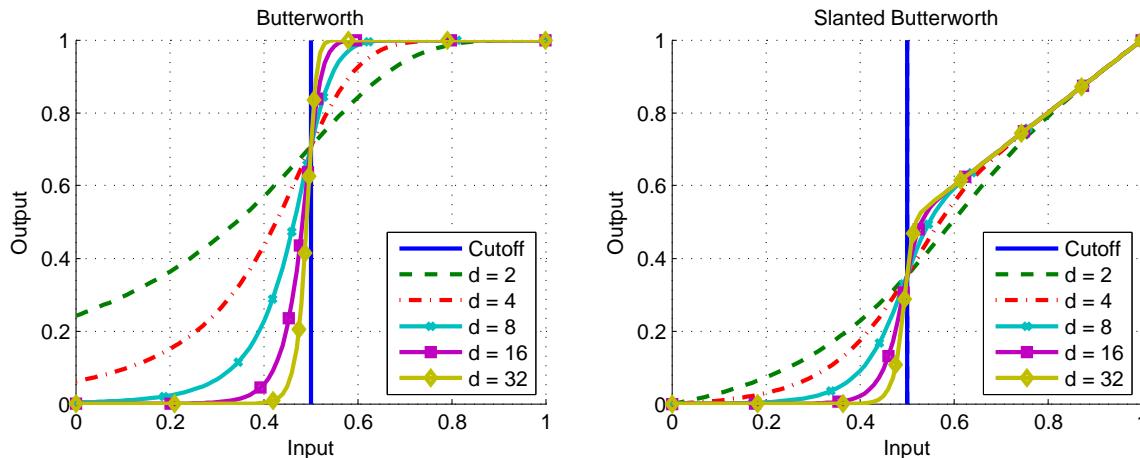


Figure 1: Plots of the Butterworth and Slanted Butterworth functions, for cutoff set at 0.5 and different filter orders.

4 Computing low-rank approximation based on Chebyshev polynomials

The evaluation of the matrix function $f_{\omega,d}^{sb}(A)$ for a large matrix A can be challenging. While the function essentially operates on the eigenvalues of the NLM operator A , the spectral decomposition of A is assumed to be unavailable due to computational reasons. Furthermore, evaluating the function $f_{\omega,d}^{sb}$ directly according to its definition (4) requires computing the square root of a matrix. Evaluating the square root of a matrix is prohibitively computationally expensive, see e.g. [9], and thus, a direct evaluation of $f_{\omega,d}^{sb}$ is also infeasible. In addition, an important observation is that one does not need the resulting matrix $f_{\omega,d}^{sb}(A)$ but only the vector $\hat{x} = f_{\omega,d}^{sb}(A)y$, as can be seen from (1).

4.1 Evaluating the SB function based on Chebyshev polynomials

The Chebyshev polynomials of the first kind of degree n are defined as

$$T_n(x) = \cos(n \arccos(x)), \quad x \in [-1, 1], \quad n = 0, 1, 2, \dots$$

These polynomials satisfy the three term recursion

$$T_n(x) = 2xT_{n-1} - T_{n-2}, \quad n = 2, 3, \dots \quad (5)$$

with $T_0(x) = 1$ and $T_1(x) = x$, and form an orthogonal basis for $L_2([-1, 1])$ with the inner product

$$\langle f, g \rangle_T = \frac{2}{\pi} \int_{-1}^1 \frac{f(t)g(t)}{\sqrt{1-t^2}} dt. \quad (6)$$

The Chebyshev expansion for any $f \in L_2([-1, 1])$ is thus given by

$$f(x) = \sum_{j=0}^{\infty} \alpha_j T_j(x), \quad \alpha_0 = \frac{1}{2} \langle f, T_0 \rangle_T, \quad \alpha_n = \langle f, T_n \rangle_T, \quad n \in \mathbb{N}, \quad (7)$$

where the equality above holds in the L_2 sense for any $f \in L_2([-1, 1])$, and becomes pointwise equality under additional regularity assumptions on f .

We will evaluate the function $f_{\omega,d}^{sb} : [0, 1] \rightarrow \mathbb{R}$ by truncating the Chebyshev expansion (7), that is

$$f_{\omega,d}^{sb}(z) \approx \sum_{j=0}^{N-1} \alpha_j T_j(y),$$

where α_j are the corresponding Chebychev coefficients for $f_{\omega,d}^{sb}$, and $y = 2z - 1$ maps $f_{\omega,d}^{sb}$ from $[0, 1]$ to $[-1, 1]$, as required by the definition of α_j above.

As shown in Corollary 3, applying a Chebychev expansion to a NLM matrix A is reduced to applying it to the diagonal form of A . Thus,

$$\begin{aligned} f_{\omega,d}^{sb}(A) &= Q^{-1} f_{\omega,d}^{sb}(\Lambda) Q \\ &= Q^{-1} \left(\sum_{j=0}^{\infty} \alpha_j T_j(2\Lambda - I) \right) Q \\ &= Q^{-1} \left(\sum_{j=0}^{\infty} \text{diag}(\alpha_j T_j(2\lambda_1 - 1), \dots, \alpha_j T_j(2\lambda_n - 1)) \right) Q \end{aligned} \quad (8)$$

where the second equality holds for any A since $f_{\omega,d}^{sb}$ is a continuous function [18, Chapter 8]. In other words, one can see that the coefficients $\{\alpha_j\}$ required to evaluate the matrix function $f_{\omega,d}^{sb}(A)$ of (4) are the same as those required to evaluate the scalar function $f_{\omega,d}^{sb}(x)$ of (3).

For square matrices, substituting a matrix inside a polynomial is well-defined (see e.g. [8, Chapter 11]) and so given a truncation parameter $N \in \mathbb{N}$, the matrix SB function (4) is approximated by

$$f_{\omega,d}^{sb}(A) \approx S_N(f_{\omega,d}^{sb}, A) = \sum_{j=0}^N \alpha_j T_j(2A - I). \quad (9)$$

The common practice for calculating Chebyshev expansions is by using a Gauss quadrature formula for (6) combined with the discrete orthogonality of the Chebyshev polynomials. Explicitly, the coefficients α_j are calculated by (7) combined with

$$\langle f, T_j \rangle_T = \frac{1}{N+1} \sum_{k=1}^{N+1} f(x_k) T_j(x_k), \quad x_k = \cos \left(\frac{\pi(k - \frac{1}{2})}{N} \right), \quad j = 0, \dots, N. \quad (10)$$

For more details see [6, Chapter 5.8].

Having obtained the Chebyshev expansion coefficients $\{\alpha_j\}_{j=0}^N$ of (9), we turn to show how to efficiently evaluate $\hat{x} = S_N(f_{\omega,d}^{sb}, A)y$. We do the latter by using a variant of Clenshaw's algorithm (see e.g., [6, Chapter 5.5]), which is based on the three term recursion (5), as described in [6, p. 193]. This algorithm is adapted to matrices using the fact that each polynomial consists only of powers of A , which means that any two matrix polynomials commute. In addition, we exploit the fact that we need to compute only the vector \hat{x} and not the matrix $S_N(f_{\omega,d}^{sb}, A)$ itself, which has much larger dimensions. The pseudocode for evaluating $S_N(f_{\omega,d}^{sb}, A)y$ for some vector y is given in Algorithm 2 in Appendix B. Note that this algorithm does not require any matrix-matrix operations, and thus feasible even for large matrix sizes.

The complete denoising algorithm, which computes a denoised image based on $S_N(f_{\omega,d}^{sb}, A)$ for a NLM operator A , is presented in Algorithm 3 in Appendix B.

4.2 Error bound for NLM operators

We study the approximation error of the truncated Chebyshev expansion for the case of matrix functions and for NLM operators, which are diagonalizable, non-symmetric matrices. The use of truncated Chebyshev expansions for matrices (9) and their approximation order has already been studied in the context of solutions for partial differential equations [17, 18]. However, most results assume that the approximated function is analytic, and so not applicable in our case.

In this section we use the following notation. Denote by $\|X\|_F = \sqrt{\text{tr}(XX^T)}$ the Frobenius norm of X and by $\|X\|$ its spectral norm (or the induced 2-norm), that is the largest singular value of X . In addition, denote by $\kappa(X)$ the condition number of an invertible matrix X (with respect to the spectral norm), that is $\kappa(X) = \|X^{-1}\| \|X\|$.

Recall that if A is a NLM operator, then A can be decomposed as $A = Q^{-1} \Lambda Q$ (see Lemma 2). In addition, A is conjugated to a symmetric matrix via $D^{-\frac{1}{2}}$ with D being the diagonal matrix of Definition 1. Therefore, $Q = D^{-\frac{1}{2}} O$, where O is an orthogonal matrix.

The next theorem presents the main error bound.

Theorem 7. Let $f \in C^{m+1}([-1, 1])$ and let A be an $n \times n$ NLM operator. Denote by

$$e_N(f)(x) = f(x) - S_N(f, x)$$

the approximation error of the truncated Chebychev expansion of degree N . Then,

$$\|e_N(f)(A)\| \leq C \frac{1}{(N-m)^m} \kappa(D^{1/2}),$$

where $C = \frac{2}{\pi m} \|f^{(m+1)}\|_T$ is a constant that depends on the $m+1$ derivative of f but independent of n and A , with

$$\|f^{(m+1)}\|_T = \int_{-1}^1 \frac{|f^{(m+1)}(t)|}{\sqrt{1-t^2}} dt.$$

Proof. Since A is diagonalizable, and similarly to Corollary 3 and (8),

$$e_N(f)(A) = f(A) - \sum_{j=0}^N \alpha_j T_j(A) = Q^{-1} (f(\Lambda) - \sum_{j=0}^N \alpha_j T_j(\Lambda)) Q = Q^{-1} E_\Lambda Q,$$

where E_Λ is the diagonal matrix $E_\Lambda = f(\Lambda) - \sum_{j=0}^N \alpha_j T_j(\Lambda)$. For all submultiplicative norms, including the spectral norm, we have that

$$\|e_N(f)(A)\| \leq \|Q^{-1}\| \|E_\Lambda\| \|Q\|. \quad (11)$$

By the Chebychev approximation bound [19, Theorem 4.3] for scalar functions,

$$\|(E_\Lambda)_{ii}\| \leq \frac{C}{(N-m)^m}, \quad 1 \leq i \leq n,$$

where $C = \frac{2}{\pi m} \|f^{(m+1)}\|_T$. The spectral norm of a diagonal matrix is given by its element on the diagonal having maximal absolute value and thus

$$\|E_\Lambda\| \leq \frac{C}{(N-m)^m}. \quad (12)$$

Now, it remains to find a bound on $\|Q^{-1}\| \|Q\|$. However,

$$\|Q^{-1}\| \|Q\| = \|O^* D^{1/2}\| \|D^{-1/2} O\|. \quad (13)$$

Using again the submultiplicativity property we get

$$\|Q^{-1}\| \|Q\| \leq \|O\| \|O^*\| \|D^{1/2}\| \|D^{-1/2}\| = \kappa(O) \kappa(D^{1/2}), \quad (14)$$

where in the last equality we have used the orthogonality of O . By the same property, we have that $\|O\| = 1$ and $\kappa(O) = 1$. Combining the latter with (11), (12), and (14) concludes the proof. \square

The proof of Theorem 7 holds with minor adjustments to various submultiplicative norms. However, for one particular norm a good bound can be achieved directly, as explained in the following remark.

Remark 2. *The error bound of the truncated Chebyshev expansion for matrices, as appears in Theorem 7, can be also expressed in terms of the Frobenius norm since $\|X\|_F \leq \sqrt{n}\|X\|$. Therefore, we immediately get*

$$\|e_N(f)(A)\|_F \leq C \frac{\sqrt{n}}{(N-m)^m} \kappa_F(D^{1/2}),$$

where $\kappa_F(D^{1/2}) = \left\|D^{1/2}\right\|_F \left\|D^{-1/2}\right\|_F$.

The bound of Theorem 7 indicates that the approximation error decays as $\frac{1}{N^m}$, where m is related to the smoothness class of the function f , and N is the number of terms in the truncated Chebyshev expansion (9). However, there are two additional factors that appear on top of this decay rate. The first one is the constant C , governed by $\left\|f^{(m+1)}\right\|_T$. This constant can be large for a function whose $m+1$ derivative has large magnitude. The second one is the condition number of $D^{1/2}$. This condition number can be easily calculated numerically since D is a diagonal matrix and thus

$$\kappa(D^{1/2}) = \frac{\max_i \sqrt{D_{ii}}}{\min_i \sqrt{D_{ii}}}. \quad (15)$$

Similarly, for the Frobenius norm we have

$$\kappa_F(D^{1/2}) = \sqrt{\frac{\sum_i D_{ii}}{\sum_i D_{ii}^{-1}}}. \quad (16)$$

Moreover, we can bound (15) and (16) a priori, as seen next.

Lemma 8. *Let D be the diagonal matrix from Definition 1. Then,*

$$\kappa(D^{1/2}) \leq \sqrt{n} \quad \text{and} \quad \kappa_F(D^{1/2}) \leq n.$$

Proof. K_h in Definition 1 is a Gaussian kernel and so $W_{ii} = 1$ and $0 \leq W_{ij} \leq 1$. Thus, $1 \leq D_{ii} \leq n$ and $1 \leq D_{ii}^{1/2} \leq \sqrt{n}$. Therefore, by (15) we have $\kappa_2(D^{1/2}) \leq \sqrt{n}$.




Figure 2: Relative approximation errors for the truncated Chebyshev expansion of $f_{\omega,d}^{sw}$ with a fixed ω and different values of d .

For the Frobenius norm, it follows that $\frac{1}{n} \leq D_{ii}^{-1} \leq 1$ and thus $\sum_i D_{ii} \leq n^2$ and $\sum_i D_{ii}^{-1} \geq 1$. Thus, by (16), $\kappa_F(D^{1/2}) \leq n$. \square

Equipped with Lemma 8, we conclude that

Corollary 9. *In the notation of Theorem 7 we have*

$$\|e_N(f)(A)\| \leq C \frac{\sqrt{n}}{(N-m)^m}.$$

Note that by Remark 2, a bound on the Frobenius norm of the approximation error $e_N(f)(A)$ is given by

$$\|e_N(f)(A)\|_F \leq C \frac{n^{1.5}}{(N-m)^m}.$$

To conclude the above discussion, we evaluate the approximation error numerically, depicted in Figure 2, where we use 50 random, positive diagonalizable, and non-symmetric matrices, whose rank equals to 1000. We average the relative approximation errors defined as $\|E_N(f_{\omega,d}^{sw}, A)\| / \|f_{\omega,d}^{sw}(A)\|$ over all 50 matrices, where the truncated Chebyshev series is evaluated by Algorithm 2. The cutoff parameter ω of $f_{\omega,d}^{sw}$ has been set to $\omega = 0.7$. The plotted curves correspond to the orders 4, 8 and 16 of the function $f_{\omega,d}^{sw}$. The results clearly show that as the derivative grows, the error rate increases, since as d gets larger so are the derivatives of $f_{\omega,d}^{sw}$.

5 Numerical experiments

The advantage of improving the NLM operator by using its low-rank approximation has already been argued in Section 3. In this section we demonstrate this advantage by numerical examples, and in addition, study the effect of the two main parameters of our method – the kernel width h from Definition 1 and the filter cutoff ω from (4). As a reference, we use the naive approach for denoising a NLM operator, by computing its truncated eigenvalues decomposition (Algorithm 1).

The numerical experiments are performed on real images corrupted by synthetically added white Gaussian noise (as described in Section 2) of varying levels. The clean test images of size 120×120 pixels are shown in Figure 3.

The denoising performance in all of the following experiments is measured by the peak signal-to-noise ratio (PSNR), which is given for images with values in the range $0 - 255$ by

$$\text{PSNR}(I_1, I_2) = 20 \cdot \log_{10} \left(\frac{255}{\sqrt{\sum_{i,j} |I_1(i, j) - I_2(i, j)|^2}} \right). \quad (17)$$

In this metric, two similar images, for example an image and a good approximation of it, get high PSNR value, while two different images have low PSNR value.

In addition to the PSNR, which measures the difference between two images, we use the signal-to-noise ratio (SNR) to measure the noise level in a given (single) image. SNR is given by


$$\text{SNR}(I_c, \sigma_{noise}) = \frac{\sigma_{I_c}}{\sigma_{noise}},$$

where I_c is the clean image, σ_{I_c} is the standard deviation of the intensity values of I_c and σ_{noise} is the standard deviation of the noise added to I_c .

The remaining of this section is organized as follows. First we examine the gain in using low-rank approximation, for different noise levels and kernel widths. Then, we compare between the naive approach of truncated spectral decomposition and our method with respect to the cutoff point, namely, the effect of different low-rank values on these two methods. We then proceed to comparing our method to a few popular denoising methods. Finally, we explore the performance of our method when combined in a state-of-the-art denoising scheme.



(a) Mandril



(b) Barbara



(c) Roof



(d) Lena



(e) Clown



(f) Couple



(g) Pentagon



(h) Woman_Blonde



(i) Man



(j) Lake



(k) Hill



(l) House



(m) Boat



(n) Man_2



(o) Couple_2

Figure 3: Test images

5.1 Effectiveness of low-rank NLM operators

We explore the improvement achieved by low-rank approximations of the NLM operator compared to the original NLM operator, as a function of the noise level (SNR) of the image and the kernel width of the operator.

For this experiment we used four of the test images given in Figure 3, resized to 60×60 pixels by bi-cubic interpolation. We computed NLM operators for a range of kernel widths with patch size $p = 5$. For each operator corresponding to a given kernel width, we constructed several low-rank approximations of it using the method of eigenvalues truncation, as given in Algorithm 1 in Appendix B. The low-rank values are given in the set

$$K = \{1, 5, 10, 15, 20, 25, 50, 75, 100, 125, 150, 175, 200, \\ 225, 250, 275, 300, 400, 600, 1200, 2000, 3000, 3600\}.$$

In Figure 4, we present three charts for every image, corresponding to three SNR levels (for 4 out of the 12 images of Figure 3). Each chart shows the PSNR between the clean image and its denoised version, as a function of the kernel width, for two operators. The two operators are the NLM operator and the best performing low-rank approximation, that is

$$\arg \max_{k_i \in K} \{ \text{PSNR}(\text{NLM-Eig}(I_c, p, k_i), I) \}, \quad (18)$$

where I and I_c are the clean and corrupted images, respectively.

From the results of this experiment we can see that the performance gap between the optimal low-rank operator (18) and the NLM operator can be very large, in favor of the low-rank approximation, for kernel widths that are much smaller than the ideal width. This advantage diminishes when the kernel width approaches its best performing value. Note that as the level of noise increases (the SNR value decreases) this optimal width value spreads to a broader range of values. In addition, one observes that the performance gain of the low-rank operator naturally diminishes as the SNR increases since less noise is involved.

5.2 The cutoff point

We investigate the effect of the cutoff point on the performance of two denoising methods. For the naive method of truncated eigenvalues (Algorithm 1), the cutoff point corresponds to the number of leading eigenvalues preserved in the low-rank operator. For our method based upon SB function (Algorithm 3), the cutoff point corresponds to the filter’s cutoff parameter ω , which is a value between zero and one (Section 3.2). In the experiment of this section, we measure the “PSNR gain”, which is the difference between the PSNR of the low-rank

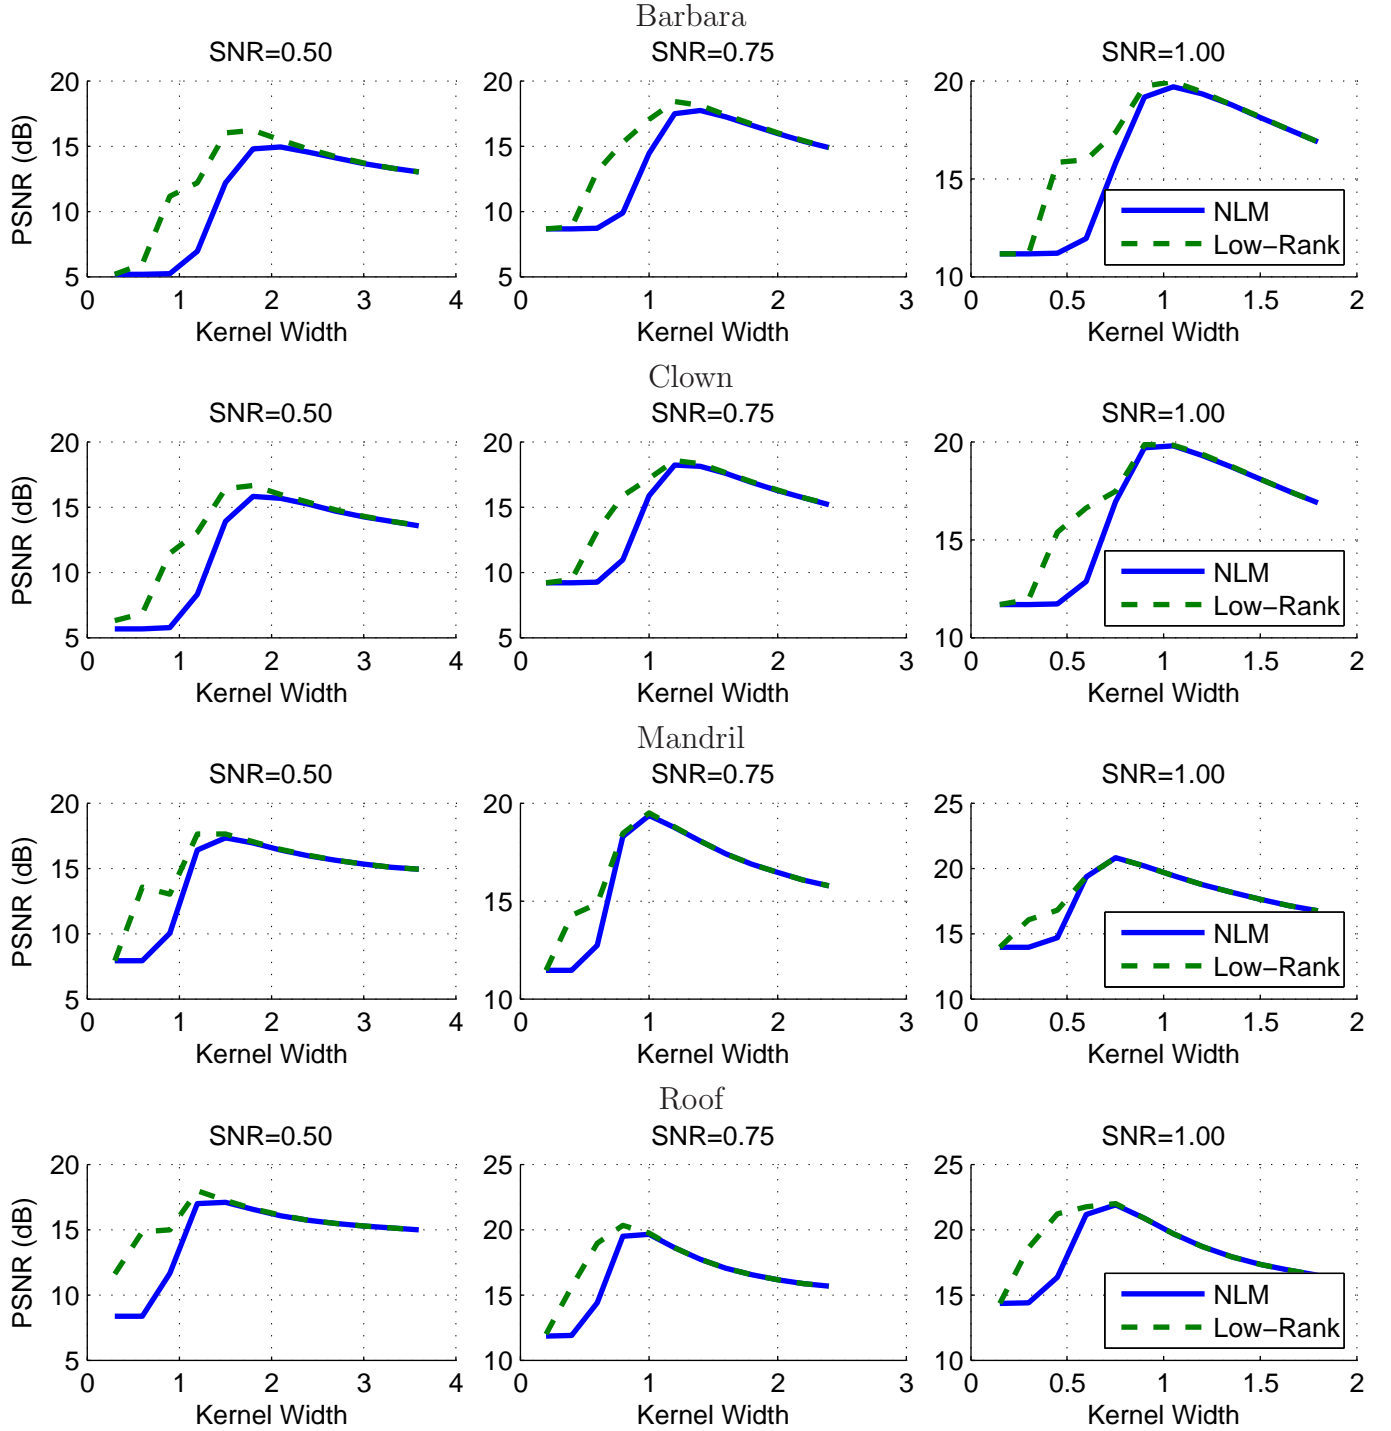


Figure 4: Using a low-rank operator versus the NLM operator: the PSNR between the clean image and its denoised version as a function of the kernel width, for different SNR (noise) levels. The original images are given in Figure 3.

operator and that of the original NLM operator it has been created from, as a function of the cutoff parameter.

The original NLM operators are constructed with a patch size $p = 5$ and kernel widths 1.5, 1.0 and 0.5 for the different SNR values 0.5, 0.75 and 1, respectively. These values were chosen based on the experiments in Section 5.1, as the kernel widths which result in the highest PSNR, on average. For our method based on the SB functions we use $N = 150$ coefficients in the truncated Chebyshev expansion (9). In addition, we select $d = 15$ as the order of the filter (4) for the higher level of noise, that is SNR of 0.5, while taking $d = 4$ for the other two SNRs of 0.75 and 1.

Figures 5 and 6 show the results of the experiment, per image and noise level using eigenvalues truncation and our method, respectively. One can observe that for both methods, the best-performing cutoff, given the SNR, is far from being the same for all images. For cutoffs that are very low (resulting in extremely low-rank operators), the PSNR gain may be negative. For cutoffs that are very high, the PSNR of the low rank operator converges to the PSNR achieved by the original NLM operator, since the modified operator itself in both methods converges to the original NLM operator.

For a different perspective, we show in Figure 7 the mean improvement per noise level for the two methods. Indeed, there is a range of cutoffs that yields an improvement on the dataset for both methods.

5.3 Comparison with other algorithms

In Figure 7 we show the optimal average cutoff value for each noise level, averaged over the entire set of test images. In this section, we compare the results achieved using these cutoffs to the results of other denoising algorithms.

The denoising algorithms we compare are: NLM, truncated eigenvalues NLM (NLM-Eig), our NLM based on the SB function (NLM-SB), NLM-SB2 (two stage denoising scheme to be described in Section 5.4 below), K-SVD [5], shape-adaptive DCT transform (SA-DCT) [7], and block-matching and 3D filtering (BM3D) [4]. The latter three have implementations which are publicly available. The parameters used for these algorithms are the default parameters given by the respective authors in their implementations.

Tables 1 and 2 contain PSNR comparisons of the above algorithms for the images of Figure 3, including a summary line consisting of the average PSNR over the full set of images.

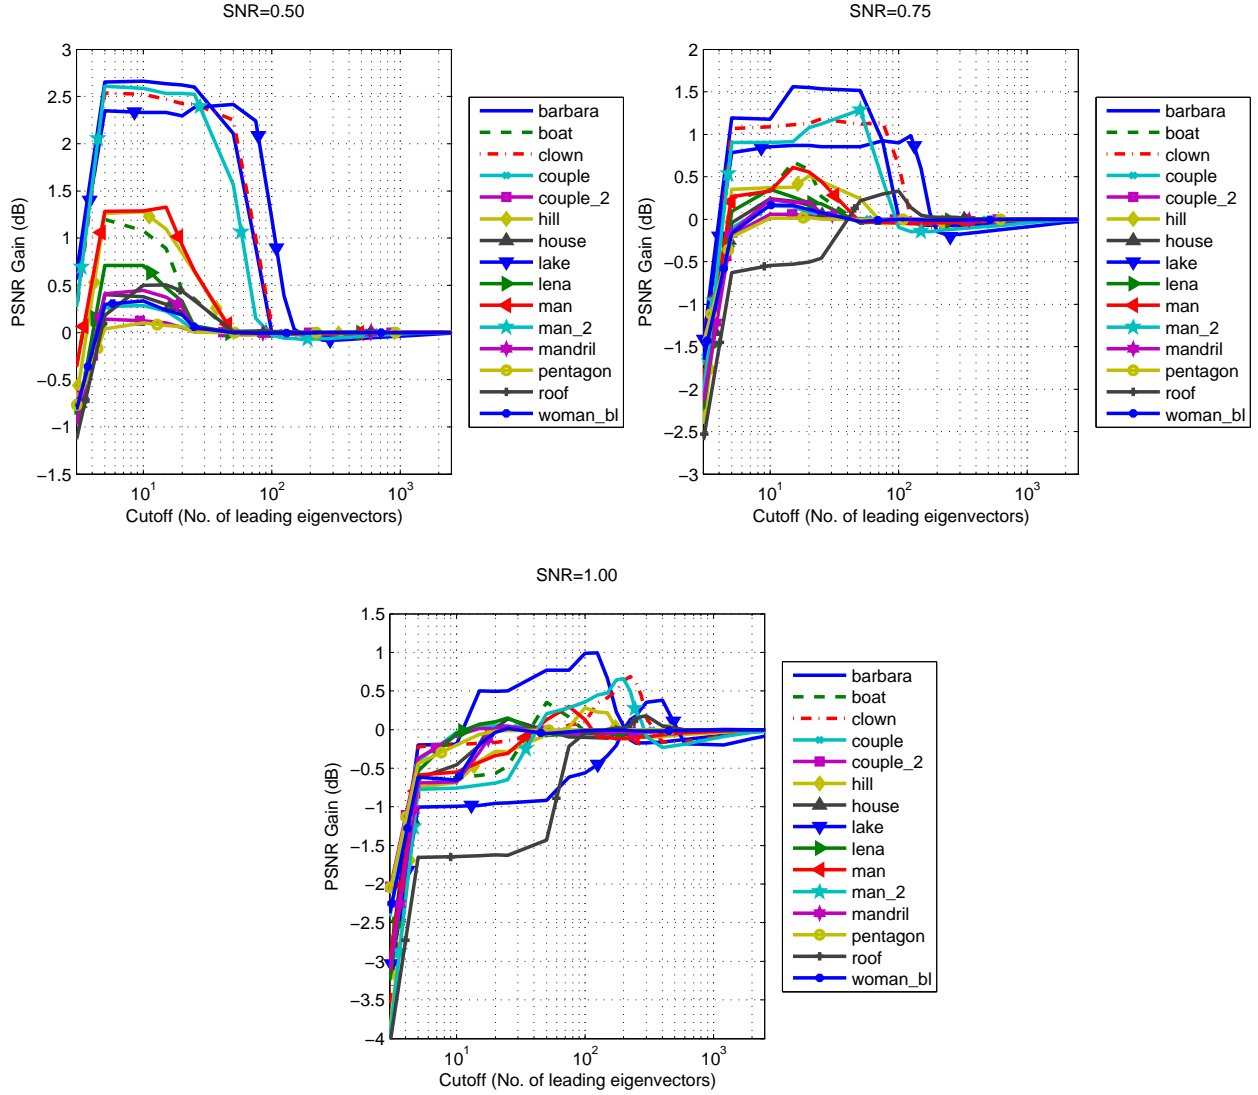


Figure 5: The differences between the PSNR of the low-rank and the original NLM operator (PSNR gain), based on truncated eigenvalues, as a function of the cutoff point (number of preserved leading eigenvalues).

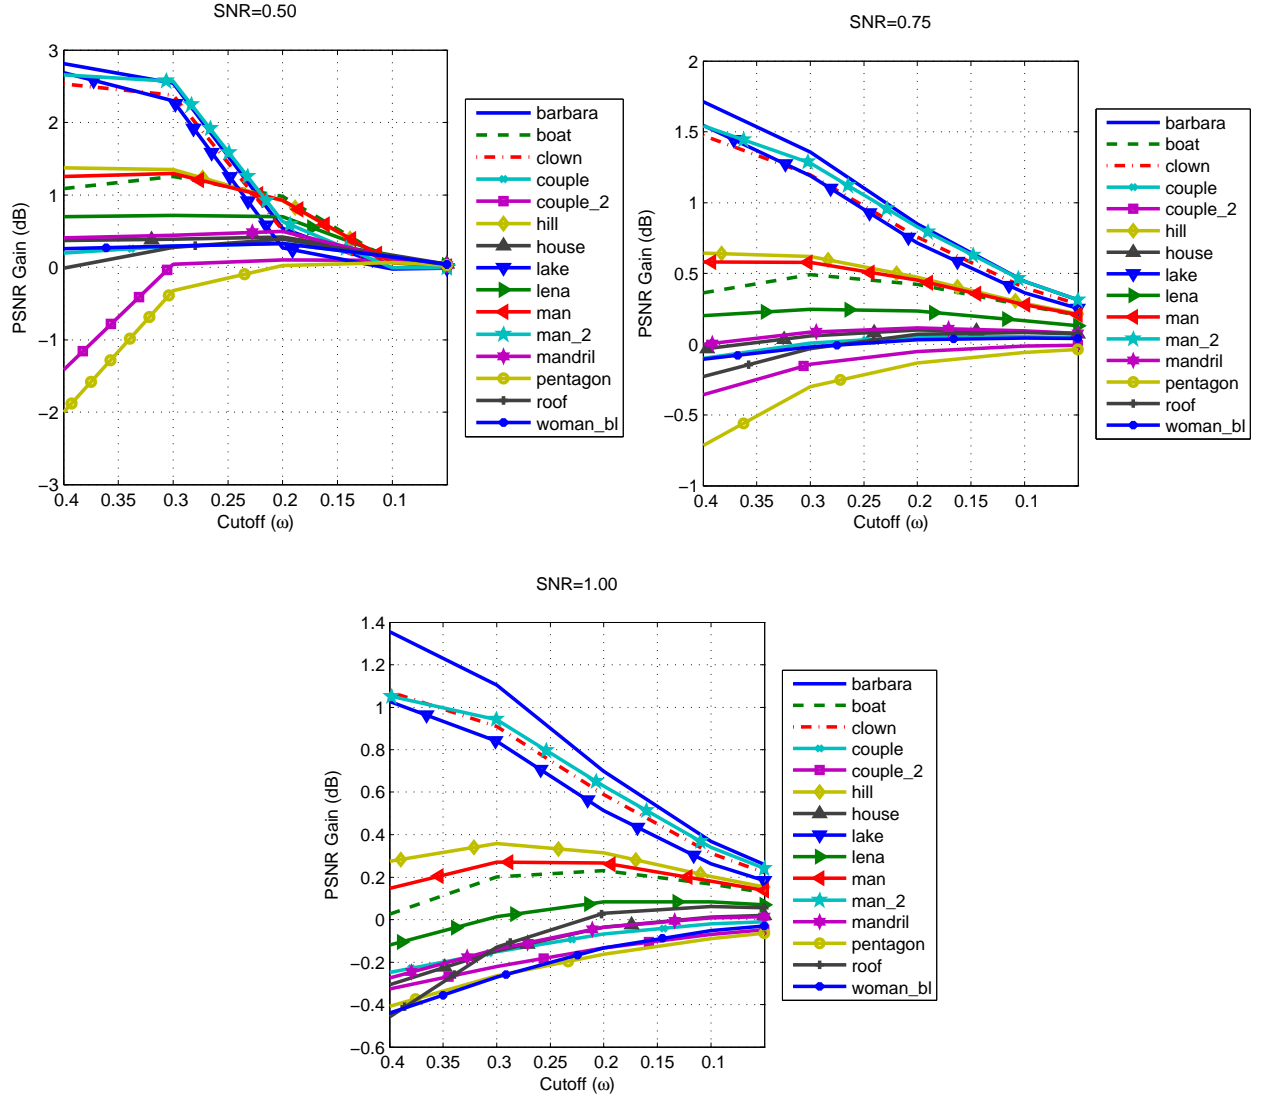


Figure 6: The differences between the PSNR of the low-rank and the original NLM operator (PSNR gain), based on the SB function, as a function of the cutoff point (the parameter ω).




Figure 7: The average PSNR improvement over the full set of test images given in Figure 3 for three levels of noise. On the left, the results of the method of truncated eigenvalues (Algorithm 1). On the right, the results of our method, based on the SB function (Algorithm 3).

Image	NLM	NLM-Eig	NLM-SB	NLM-SB2	K-SVD	SA-DCT	BM3D
barbara	15.47	18.13	18.14	19.37	20.10	21.09	20.88
boat	16.98	18.06	18.20	19.13	18.26	20.40	20.81
clown	15.48	18.00	17.94	18.70	18.14	19.66	19.85
couple	17.97	18.25	18.25	18.69	20.03	21.00	21.24
couple_2	18.62	18.74	18.70	18.91	19.79	22.20	22.58
hill	16.56	17.84	17.86	18.86	18.54	20.09	20.30
house	18.15	18.52	18.52	19.27	18.95	22.38	23.44
lake	14.81	17.14	17.24	18.36	17.75	18.98	18.83
lena	17.52	18.23	18.21	19.01	18.90	20.61	20.99
man	16.57	17.86	17.82	18.83	17.88	19.56	19.75
man_2	15.82	18.41	18.46	19.51	19.23	20.37	20.49
mandril	17.30	17.75	17.72	18.37	17.84	20.05	20.27
pentagon	18.07	18.17	18.04	17.96	15.56	20.55	21.27
roof	17.33	17.83	17.52	18.15	17.19	20.64	21.80
woman_blonde	18.02	18.36	18.31	18.85	18.00	20.77	21.05
Average	16.98	18.09	18.06	18.80	18.41	20.56	20.90

Table 1: PSNR values of the tested algorithms for noise level corresponding to SNR=0.5.

Image	NLM	NLM-Eig	NLM-SB	NLM-SB2	K-SVD	SA-DCT	BM3D
barbara	18.83	20.39	20.65	21.27	20.87	22.49	22.90
boat	19.68	20.35	20.15	21.13	18.64	21.92	22.22
clown	18.66	19.77	20.16	20.37	19.00	21.32	21.62
couple	20.06	20.21	19.99	21.86	19.93	22.52	22.85
couple_2	20.48	20.54	20.33	22.38	20.88	23.78	24.15
hill	19.41	19.79	20.00	21.25	18.97	21.72	21.94
house	20.78	20.99	20.76	22.74	20.75	24.67	25.58
lake	18.05	18.92	19.55	18.70	18.22	20.48	20.49
lena	20.02	20.28	20.14	21.35	19.43	22.26	22.57
man	19.45	20.06	19.97	20.66	18.28	21.08	21.30
man_2	19.13	20.04	20.74	21.06	19.44	21.98	22.16
mandril	19.50	19.70	19.44	20.76	19.51	21.15	21.39
pentagon	19.51	19.52	19.31	20.82	19.05	22.31	22.90
roof	19.84	19.30	19.53	20.71	19.98	22.12	23.73
woman_blonde	20.31	20.47	20.15	21.26	19.80	22.55	22.87
Average	19.58	20.02	20.06	21.09	19.52	22.16	22.58

Table 2: PSNR values of the tested algorithms for noise level corresponding to SNR=0.75.

5.4 Two stages denoising schemes

Meyer and Shen [11] have proposed to compute a second NL-means operator from the denoised image and apply it again. Except for introducing two stages, there are a few additional differences between their algorithm and the NL-means algorithm, given in Definition 1.

First, it contains a k -nearest neighbours search for constructing the denoising operators. Instead of constructing the operator $A = D^{-1}W$, where the elements of W are given as in Definition 1, they construct W as $W = 0.5(Z + Z')$, where the matrix Z contains k nonzero elements in row i , which are the k largest values of the set

$$\left\{ K_h(v(N_i^Y) - v(N_j^Y)) \mid 1 \leq j \leq n \right\}.$$

Second, the way that the denoising operator is applied is different than in the NLM algorithm. Instead of computing the denoised image as $\hat{x} = Ay$, where each pixel is denoised using only the patch of which it is the center, they average the values of the pixel from all patches in which it is included.

For comparison purposes, we have modified Meyer’s scheme to use our low-rank operator, based on the SB function. The resulting algorithm is henceforth referred to as SB-Meyer. In order to use $f_{\omega,d}^{sb}$ in this scheme, we had to modify the way the denoising operator is computed; In contrary to the NLM operator, the operator in Meyer’s scheme is not semi-positive definite (not all eigenvalues are non-negative). This is due to using k -nearest neighbours to construct

it. Since our framework, as given in Section 3, requires a non-negative spectrum, we have approximated the denoising operator in Meyer’s scheme with a new semi-positive definite matrix, by shifting its spectrum such that the largest negative eigenvalue becomes zero. Then, we normalized its rows such that they will sum to unity. This normalization transforms the largest positive eigenvalue back to unity (see also the proof of Lemma 2 in Appendix A).

The original algorithm of [11] uses parameters given by the authors in their source code, which were tuned only for SNR level of 1. Thus, we report in Table 3 its results (in the column “Meyer”) and also those of SB-Meyer only for that noise level. The comparison is over the full set of images (as given in Figure 3), where the parameters of SB-Meyer are given in Table 4. In addition to SB-Meyer, we have implemented a simpler two-stage scheme employing our operator based on the SB function, which is given in Algorithm 4 and is referred to as NLM-SB2. Its parameters are given in Table 5. From Table 3 we see that using the method of the current paper improves the resulting PSNR (on average). Using a two stage scheme further improves the resulting PSNR.

Image	NLM	NLM-Eig	NLM-SB	NLM-SB2	SB-Meyer	Meyer	K-SVD	SA-DCT	BM3D
barbara	20.84	21.83	21.94	22.11	23.18	20.30	21.67	23.61	24.07
boat	21.24	21.16	21.44	22.01	22.72	21.33	21.40	23.08	23.22
clown	20.51	20.85	21.42	21.08	22.23	20.14	20.40	22.54	22.85
couple	21.66	21.65	21.51	22.54	23.50	23.06	22.27	23.86	24.17
couple_2	22.03	22.03	21.81	23.17	24.46	24.61	22.20	25.17	25.37
hill	21.27	21.50	21.63	22.20	22.57	21.21	21.38	23.00	23.15
house	22.83	22.72	22.68	23.89	24.82	26.19	24.10	26.31	27.18
lake	19.73	19.28	20.57	19.24	21.42	20.78	19.77	21.66	21.66
lena	21.73	21.68	21.75	22.39	23.12	22.44	21.63	23.51	23.74
man	21.18	21.09	21.45	21.70	22.04	21.16	19.94	22.32	22.41
man_2	20.94	21.38	21.88	21.47	22.08	20.01	20.02	23.14	23.29
mandril	21.08	21.03	20.94	21.45	22.37	21.50	20.61	22.11	22.24
pentagon	20.84	20.84	20.57	21.61	23.50	24.22	21.62	23.67	24.12
roof	21.97	21.93	21.84	22.66	22.57	24.42	21.64	23.46	25.04
woman_blonde	21.87	21.86	21.60	21.89	23.18	23.64	23.17	24.03	24.22
Average	21.31	21.39	21.53	21.96	22.92	22.33	21.45	23.43	23.78

Table 3: Algorithm comparison for SNR=1

	v_1	v_2	p_1	p_2	h_1	h_2	ω_1	ω_2	d_1	d_2	γ	k
SNR=1	200	200	7	3	1	1	0.6	0.4	50	16	0.2	150

Table 4: The parameters of the SB-Meyer scheme.

SNR	p	h_1	h_2	ω_1	ω_2	d_1	d_2	γ	N
0.5	5	1.5	1	0.3	0.3	50	50	0.5	150
0.75	5	1.05	0.35	0.3	0.3	15	15	0.15	150
1	5	0.5	0.3	0.3	0.3	4	4	0.15	150

Table 5: The parameters of the NLM-SB2 scheme, as given in Algorithm 4.

For a more visual perspective of the comparison, we provide Figures 8 and 9, where two images are denoised by the various tested algorithms. The first image is part of the Mandril image, whose clean version is in Figure 8a. Its denoised versions are presented in Figures 8c–8f. One can see that SB-Meyer retained the texture much better than BM3D, and it also contains less artifacts than Meyer’s original scheme (the finding regarding texture preservation is consistent with the conclusions in [11]). The second image is part of the Barbara image, shown in Figure 9a. The denoising results for the Barbara image are presented in Figures 9c–9f, where we can see that the outputs of NLM-SB and NLM-SB2 are considerably less noisy than the output of NLM while retaining nearly the same amount of details. In addition, it appears that some small image regions that were incorrectly estimated by NLM were also incorrectly estimated by NLM-SB. A few of these artifacts are visible on the Barbara image and they are marked by red rectangles in Figures 9d and 9e. This effect of a remaining noisy region was amplified by the application of a second denoising stage in NLM-SB2. We believe that these artifacts can be removed by some simple heuristic, but this is outside the scope of this work.

6 Summary

In this paper, we have investigated the idea of improving a Non-Local Means operator by manipulating its spectrum. We have shown a method to do so without computing explicitly neither the eigenvectors of the original operator nor the matrix of the modified operator.

Our method operates by applying a filtering function to the original Non-Local Means operator. To that end, we have derived sufficient conditions for such filtering functions and an efficient procedure for its application.

In this work we also show the connection between spectral shaping of the operator and




Figure 8: Denoised examples taken from the Mandril image with SNR= 1.

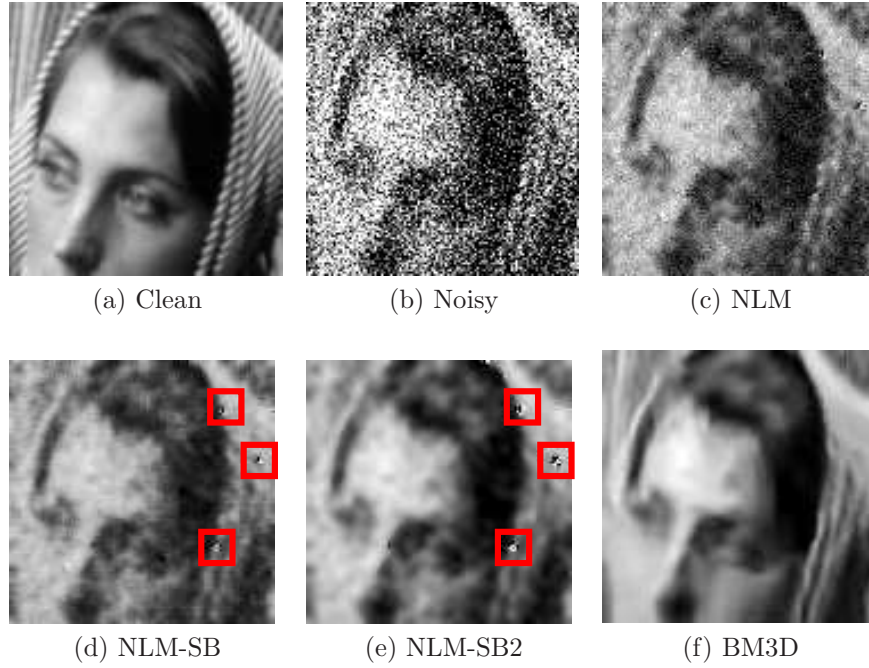


Figure 9: Denoised examples taken from the Barbara image with SNR= 0.75. Remaining noise artifacts are marked with red rectangles.

the application of a matrix function on the operator. In the implementation of our approach, we demonstrate the well-known efficiency of the Chebyshev polynomials for matrix functions. Moreover, a bound on the approximation error of the truncated Chebyshev expansion for the class of Non-Local Means matrices is proved.

We present numerical experiments from which we learn a few important observations about the improvement that can be gained by a low-rank approximation of a Non-Local Means operator. First, the improvement depends on the choice of the kernel width parameter and the noise level. Second, we find that the optimal cutoff of the spectrum of the Non-Local Means operator varies for different images and noise levels. Nevertheless, the methods of eigenvalue truncation and SB filtering, as given in Algorithms 1 and 3 respectively, both achieve a non-negligible improvement, on average, on our dataset of 15 images. It is also demonstrated that a further PSNR improvement can be achieved by running SB filtering via a two-stage scheme, as suggested in [11].

References

- [1] Antoni Buades, Bartomeu Coll, and J-M Morel. A non-local algorithm for image denoising. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 2, pages 60–65. IEEE, 2005.
- [2] Stephen Butterworth. On the theory of filter amplifiers. *Wireless Engineer*, 7:536–541, 1930.
- [3] Ronald R Coifman and David L Donoho. *Translation-invariant de-noising*. Springer New York, 1995.
- [4] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, Karen Egiazarian, et al. BM3D image denoising with shape-adaptive principal component analysis. In *SPARS’09-Signal Processing with Adaptive Sparse Structured Representations*, 2009.
- [5] Michael Elad and Michal Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *Image Processing, IEEE Transactions on*, 15(12):3736–3745, 2006.
- [6] Brian P Flannery, William H Press, Saul A Teukolsky, and William Vetterling. Numerical recipes in C. *Press Syndicate of the University of Cambridge, New York*, 1992.

- [7] Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Pointwise shape-adaptive dct for high-quality denoising and deblocking of grayscale and color images. *Image Processing, IEEE Transactions on*, 16(5):1395–1411, 2007.
- [8] Gene H Golub and Charles F Van Loan. *Matrix computations*, volume 3. JHU Press, 2012.
- [9] Nicholas J Higham. Stable iterations for the matrix square root. *Numerical Algorithms*, 15(2):227–242, 1997.
- [10] Aapo Hyvärinen, Jarmo Hurri, and Patrik O Hoyer. *Natural Image Statistics: A Probabilistic Approach to Early Computational Vision.*, volume 39. Springer, 2009.
- [11] François G Meyer and Xilin Shen. Perturbation of the eigenvectors of the graph laplacian: Application to image denoising. *Applied and Computational Harmonic Analysis*, 36(2):326–334, 2014.
- [12] Giovanni Motta, Erik Ordentlich, Ignacio Ramirez, Gadiel Seroussi, and Marcelo J Weinberger. The idude framework for grayscale image denoising. *Image Processing, IEEE Transactions on*, 20(1):1–21, 2011.
- [13] Pietro Perona and Jitendra Malik. Scale-space and edge detection using anisotropic diffusion. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 12(7):629–639, 1990.
- [14] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1):259–268, 1992.
- [15] Amit Singer, Yoel Shkolnisky, and Boaz Nadler. Diffusion interpretation of nonlocal neighborhood filters for signal denoising. *SIAM Journal on Imaging Sciences*, 2(1):118–139, 2009.
- [16] Richard Szeliski, Ramin Zabih, Daniel Scharstein, Olga Veksler, Vladimir Kolmogorov, Aseem Agarwala, Marshall Tappen, and Carsten Rother. A comparative study of energy minimization methods for markov random fields with smoothness-based priors. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 30(6):1068–1080, 2008.
- [17] Hillel Tal-Ezer. Polynomial approximation of functions of matrices and applications. *Journal of Scientific Computing*, 4(1):25–60, 1989.
- [18] Lloyd N Trefethen. *Spectral methods in MATLAB*, volume 10. SIAM, 2000.

- [19] Lloyd N Trefethen. Is Gauss quadrature better than Clenshaw-Curtis? *SIAM review*, 50(1):67–87, 2008.
- [20] Maria Zontak and Michal Irani. Internal statistics of a single natural image. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 977–984. IEEE, 2011.

Appendix

A Proof of Lemma 2

Lemma 2 summarizes known properties of NLM operators. We provide its proof for the self-containedness of the paper.

Proof. The NLM operator of Definition 1 is in general not symmetric, but it is conjugated via $D^{\frac{1}{2}}$ to the symmetric matrix $S = D^{-\frac{1}{2}}WD^{-\frac{1}{2}}$, namely $A = D^{-1/2}SD^{1/2}$. Therefore, the NLM operator is diagonalizable and has the same eigenvalues as S . On other hand, K_h is the Gaussian kernel function, which implies that W is a symmetric positive definite matrix. Since S is obtained from W by multiplication by $D^{-1/2}$ on both sides, by Sylvester’s law of inertia, the eigenvalues of S are positive as well. Thus, since S and A are conjugated, all eigenvalues of A are also positive.

Since A is element-wise positive, it follows from the Perron-Frobenius theorem that

$$\min_{1 \leq i \leq n} \sum_{j=1}^n A_{ij} \leq \left| \max_{1 \leq i \leq n} (\lambda_i) \right| \leq \max_{1 \leq i \leq n} \sum_{j=1}^n A_{ij}.$$

Moreover, since for any $1 \leq i \leq n$, $\sum_j A_{ij} = 1$, we get that $\left| \max_{1 \leq i \leq n} (\lambda_i) \right| = 1$. Thus, $\lambda_1 = 1$ and from the positivity of the eigenvalues we conclude that $\lambda_i \in (0, 1]$, $1 \leq i \leq n$. \square

B Algorithms

This appendix contains the pseudocode for the algorithms described in the paper. In these algorithms we denote by $\text{CHEB-COEF}(f, N)$ the set of the first $N + 1$ Chebyshev coefficients calculated via (7) and (10).

Algorithm 1 NLM-Eig denoising scheme. Here, in the notation of Lemma 2, $\text{EIG}(A, k) = Q^{-1}\Lambda_k Q$, where $\Lambda_k = \text{diag}(\lambda_1, \dots, \lambda_k, 0, \dots, 0)$.

Input: Noisy image Y , patch size p , kernel width h , matrix approximation rank k .

Output: Denoised image \hat{X} .

```

1: procedure NLM-EIG( $Y, p, h, k$ )
2:    $A \leftarrow \text{NLM}_{p,h}(Y)$  ▷ Create a NLM operator for the image.
3:    $B \leftarrow \text{EIG}(A, k)$  ▷ Make a low-rank approximation of the NLM operator.
4:    $y \leftarrow \text{COL}(Y)$  ▷ Convert the image  $Y$  to a vector.
5:    $\hat{x} \leftarrow By$  ▷ Compute the output image using the low-rank approximation of  $A$ .
6:   return IMAGE( $\hat{x}$ ) ▷ Reshape  $\hat{x}$  as an image.
7: end procedure

```

Algorithm 2 Evaluate the product of a truncated matrix Chebyshev expansion (9) by a vector.

Input: Matrix A , vector y , vector of coefficients c of length N .

Output: The vector $\sum_{k=1}^N c_k T_k(A) y$.

```

1: procedure CLENSHAWMATVEC( $A, y, c$ )
2:    $N \leftarrow |c|$ 
3:    $T \leftarrow 2 \cdot A - I_n$ 
4:    $d \leftarrow 0$ 
5:    $dd \leftarrow 0$ 
6:   for  $i \leftarrow N$  downto 2 do
7:      $temp \leftarrow d$ 
8:      $d \leftarrow 2 \cdot T \cdot d - dd + c_i \cdot y$ 
9:      $dd \leftarrow temp$ 
10:  end for
11:  return  $T \cdot d - dd + 0.5 \cdot c_1 \cdot y$ 
12: end procedure

```

Algorithm 3 NLM-SB denoising scheme.

Input: Noisy image Y , patch size p , kernel width h , cutoff and order parameters ω and d of $f_{\omega,d}^{sb}$, number of Chebychev coefficients N .

Output: Denoised image \hat{x} .

```

1: procedure NLM-SB( $Y, p, h, \omega, d, N$ )
2:    $\{\alpha_j\} \leftarrow \text{CHEB-COEF}(f_{\omega,d}^{sb}, N)$  ▷ Compute the Chebychev coefficients for  $f_{\omega,d}^{sb}$ .
3:    $A \leftarrow \text{NLM}_{p,h}(Y)$  ▷ Create a NLM operator for the image.
4:    $x \leftarrow \text{COL}(Y)$  ▷ Convert the image  $Y$  to a vector.
5:    $\hat{x} \leftarrow \text{ClenshawMatVec}(A, \{\alpha_j\}, x)$  ▷ Algorithm 2
6:   return IMAGE( $\hat{x}$ ) ▷ Reshape  $\hat{x}$  as an image.
7: end procedure

```

Algorithm 4 NLM-SB2 two-stage denoising scheme.

Input: Noisy image Y , patch size p , kernel widths h_1 and h_2 , mixing weight $\gamma \in [0, 1]$, cutoff and order parameters ω_1, ω_2 and d_1, d_2 of $f_{\omega,d}^{sb}$, number of Chebychev coefficients N .

Output: Denoised image $\hat{x}^{(3)}$.

- 1: **procedure** NLM-SB2($Y, p, \gamma, h_1, h_2, \omega_1, \omega_2, d_1, d_2$)
 - 2: $\hat{x}^{(1)} \leftarrow \text{NLM-SB}(Y, p, h_1, \omega_1, d_1, N)$ \triangleright Algorithm 3.
 - 3: $\hat{x}^{(2)} \leftarrow (1 - \gamma)\hat{x}^{(1)} + \gamma x$ \triangleright Mix the esimated image with the original one.
 - 4: $\hat{x}^{(3)} \leftarrow \text{NLM-SB}(\text{IMAGE}(x^{(2)}), p, h_2, \omega_2, d_2, N)$ \triangleright Denoise again.
 - 5: **return** $\hat{x}^{(3)}$
 - 6: **end procedure**
-